




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : H04Q 11/04, H04L 12/56		A1	(11) International Publication Number: WO 00/36872
			(43) International Publication Date: 22 June 2000 (22.06.00)
(21) International Application Number: PCT/SE99/02319 (22) International Filing Date: 10 December 1999 (10.12.99) (30) Priority Data: 09/213,897 17 December 1998 (17.12.98) US (71) Applicant: TELEFONAKTIEBOLAGET LM ERICSSON (publ) [SE/SE]; S-126 25 Stockholm (SE). (72) Inventors: ANDERSSON, Staffan; Lessebovägen 133, S-122 42 Enskede (SE). AGNEVIK, Mikael; Diligensvägen 30, S-131 48 Nacka (SE). ENGSTRÖM, Ola; Fågelstavägen 29, S-124 33 Bandhagen (SE). PETERSEN, Lars-Göran; Hökbursvägen 5, S-147 42 Tumba (SE). (74) Agent: NORIN, Klas; Ericsson Radio Systems AB, Common Patent Dept., S-164 80 Stockholm (SE).		(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>	

(54) Title: INTERNAL ROUTING THROUGH MULTI-STAGED ATM NODE



(57) Abstract

An internal routing tag (82) is appended to a payload of a received ATM cell upon ingress of the ATM cell into a multi-stage ATM node (20). The routing tag comprises routing information for routing the payload of the received ATM cell through plural stages (22) of the multi-stage node. Preferably the routing information comprises a list of destination addresses, e.g., utopia address of physical units (26) in the multi-stage ATM node. Use of the internal routing tag streamlines connection setup, reduces the number of internal control paths required, and abviates VPI/VCI conversion at the plural stages of the multi-stage node.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PC

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

INTERNAL ROUTING THROUGH MULTI-STAGED ATM NODE

BACKGROUND

This application is related to United States Patent Application Serial No. 09/213,898, entitled "SEGMENTED PERFORMANCE MONITORING OF MULTI-STAGE ATM NODE", and incorporated by reference herein.

1. FIELD OF THE INVENTION

5 The present invention pertains to the packet technology known as Asynchronous Transfer Mode (ATM), and particularly to internal routing of traffic cells through a multi-staged ATM node.

2. RELATED ART AND OTHER CONSIDERATIONS

10 The increasing interest for high band services such as multimedia applications, video on demand, video telephone, and teleconferencing has motivated development of the Broadband Integrated Service Digital Network (B-ISDN). B-ISDN is based on a technology know as Asynchronous Transfer Mode (ATM), and offers considerable extension of telecommunications capabilities.

15 ATM is a packet-oriented transfer mode which uses asynchronous time division multiplexing techniques. Packets are called cells and traditionally have a fixed size. A traditional ATM cell comprises 53 octets, five of which form a header and forty eight of which constitute a "payload" or information portion of the cell. The header of the ATM cell includes two quantities which are used to identify a connection in an ATM network over which the cell is to travel, particularly the VPI (Virtual Path Identifier) and VCI
20 (Virtual Channel Identifier). In general, the virtual is a principal path defined between two switching nodes of the network; the virtual channel is one specific connection on the respective principal path.

At its termination points, an ATM network is connected to terminal equipment, e.g., ATM network users. Typically between ATM network termination points there are plural switching nodes, the switching nodes having ports which are connected together by physical transmission paths or links. Thus, in traveling from an originating terminal equipment to a destination terminal equipment, ATM cells forming a message may travel through several switching nodes.

A switching node has a plurality of ports, each of which can be connected by via a link circuit and a link to another node. The link circuit performs packaging of the cells according to the particular protocol in use on the link. A cell incoming to a switching node may enter the switching node at a first port and exit from a second port via a link circuit onto a link connected to another node. Each link can carry cells for plural connections, a connection being e.g., a transmission between a calling subscriber or party and a called subscriber or party.

The switching nodes each typically have several functional parts, a primary of which is a switch core. The switch core essentially functions like a cross-connect between ports of the switch. Paths internal to the switch core are selectively controlled so that particular ports of the switch are connected together to allow a message ultimately to travel from an ingress side of the switch to an egress side of the switch. The ports of the switch core are connected to interface modules which can reside on device boards. The interface modules serve to interface the switch core with one or more devices also residing on the device board, such devices being (for example) processors, an AAL or ATM termination unit, or an exchange terminal (ET). Some of these devices can be connected to external links, such as (for example) an exchange terminal (ET) connects the ATM node to another ATM node via an external link. For functions such as a conversion operation hereinafter described, the exchange terminals typically have processors (known as "board processors" or "BPs") mounted thereon.

Typically an ATM switch core and its connected device boards [upon which the interface modules and devices such as exchange terminals (ET) are mounted] reside together on one subrack of a rack of electronic components. When a connection is to be set up involving the ATM node, a connection set up operation is performed. Thereafter, while the connection is set up, ATM cells are received at the ATM node on a particular external link. When an ATM cell is received, the board processor BP on the device board which is connected to the receiving external link consults conversion tables

maintained by the board processor. From the conversion tables the board processor determines a mapping from, e.g., the external VPI/VCI values (from the header of the ATM cell received on the external link) to internal VPI/VCI values. The internal VPI/VCI is used for routing the payload of the received cell through the ATM node, and particularly through its switch core. After the cell is routed through the ATM node, a similar conversion process is performed at the device board from which the cell is to leave the ATM node. That is, another conversion process uses the internal VPI/VCI to prepare another external VPI/VCI to insert in the header of an outgoing cell as it leaves the ATM node.

10 The conversion tables utilized for the VPI/VCI conversions are communicated or updated to the board processors at connection set up from by a main processor of the node. The main processor can reside, for example, on one of the device boards of the node. An internal control path (ICP) is required to connect the main processor to each of the board processors (BPs) which perform VPI/VCI conversion operations.

15 There are envisioned larger ATM nodes comprising, for example, multiple subracks with one of the subracks acting as a main switch (e.g., a "group switch" in some parlance). For example, the subracks (each constituting a stage of the overall node) can be connected in cascading fashion to form a multi-staged ATM node. The conversion process described above can be employed in such larger modes. Consider, 20 for example, a three stage ATM node comprising three subracks, with a second of the subracks functioning as the main switch and being connected between the first and third subracks. In such three stage ATM node, the routing of a cell through the node in accordance with the conversion procedure described above requires two conversion operations at each subrack (one conversion operation upon ingress; another conversion operation upon egress), for a total of six conversions. Moreover, upon setting up a 25 connection from, e.g., a device/ET in one subrack to a device/ET in another subrack requires setting up the conversion tables for all exchange terminals (ETs) of the involved internal ET links (i.e., the links internconnecting the main switch along the path of the connection). Such connection set up involves more internal control paths 30 (ICPs) and essentially increases connection set up time by a factor of three.

What is needed, therefore, and an object of the present invention, is a technique for internally routing ATM cells through a multi-staged ATM node.

BRIEF SUMMARY OF THE INVENTION

An internal routing tag is appended to a payload of an ATM cell for routing the ATM cell through a multi-stage ATM node. The routing tag comprises routing information for routing the payload of the received ATM cell through plural stages of the multi-stage node. Preferably the routing information comprises a list of destination addresses, e.g., utopia address of physical units in the multi-stage ATM node. In an illustrated embodiment, the multi-stage ATM node has stages, each stage being situated on a subrack. Each stage or subrack of the multi-stage node comprises a switch core connected to a first set of interface units and a second set of interface units. For such
5 embodiment, the routing tag includes, as the destination addresses, physical addresses for one of the first set of interface units and one of the second set of interface units for each stage of the multi-stage node. The interface units can be, for example, switch port interface modules (SPIM).
10

A connection set up request is forwarded to a connection setup manager which preferably resides at a main processor of the multi-stage ATM node. The connection
15 setup manager responds by providing transfer lists for the connection to two tagging units BP/TUs. The two tagging units BP/TUs which receive transfer lists are on device boards connected to the two external links involved in the connection. The transfer list is a list of destination addresses to be used for switching and routing of the traffic cell
20 through the multi-stage ATM node. The transfer list received by one of the two tagging units BP/TUs is used for routing cells in one direction through the multi-stage ATM node; the transfer list received by the other of the two tagging units BP/TUs is used for routing cells in a reverse direction through the multi-stage ATM node. The transfer lists are written into conversion tables of the two tagging units BP/TUs. After connection
25 setup, when an ATM cell participating in the connection is received from an external link, a tagging unit BP/TU obtains the VPI/VCI and link identifier from the header of the incoming ATM cell. Using these parameters, the tagging unit BP/TU consults its conversion table to obtain the transfer list for the cell. Two other parameters, specifically cell size and QoS indicator, are also added to form the internal routing tag
30 or "SPAS tag".

In an illustrated embodiment, the transfer list of the SPAS tag includes six octets, each octet including, e.g., a destination address for routing through the multi-stage ATM node. As a cell is routed through the multi-stage ATM node, units which receive

the cell analyze and, where appropriate, operate on the top octet in the transfer list. These operations including, as a top octet is utilized in route to or at a unit having the destination address thereof, exchanging the destination address with a source address from whence the cell came; changing the parity of the octet from odd to even; and, 5 popping the octet so that it goes to the bottom of the transfer list. The popping of the octet leaves another octet on top of the transfer list, that another octet containing the next destination address to which the cell is to be routed.

Formats of SPAS tags are provided for accommodating ATM packets of various protocols, including an AAL2' protocol and an AAL2'' protocol, as well as for various 10 Utopia devices (e.g., eight bit and sixteen bit). For multi-stage ATM nodes having a ring or bus topology, an octet pair is included in SPAS tag for routing the cell on the ring or bus.

Use of the internal routing tag streamlines connection setup, reduces the number of internal control paths required, and obviates VPI/VCI conversion at the plural stages 15 of the multi-stage node.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, features, and advantages of the invention will be apparent from the following more particular description of preferred embodiments as illustrated in the accompanying drawings in which reference characters refer to the 20 same parts throughout the various views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention.

Fig. 1 is a schematic view of portions of a multi-stage ATM node according to an example embodiment of the invention, particularly access subracks and a main rack thereof.

25 Fig. 2 is a schematic view of an portions of the multi-stage ATM node of Fig. 1 and further including a connection setup manager, a node performance monitoring manager, and a traffic manager.

Fig. 3 is a schematic view showing an example device board upon which a switch port interface module (SPIM) of the ATM node of Fig. 1 can reside.

Fig. 4 is a diagrammatic view of an example cell having a SPAS tag applied thereto for transit through the multi-stage ATM node of Fig. 1.

Fig. 4A is a diagrammatic view of an example cell having a SPAS tag applied thereto for transit through the multi-stage ATM node of Fig. 1, the example cell having
5 an AAL2' protocol packet.

Fig. 4B is a diagrammatic view of an example cell having a SPAS tag applied thereto for transit through the multi-stage ATM node of Fig. 1, the example cell having an AAL2'' protocol packet.

Fig. 4C is a diagrammatic view of an example cell having a SPAS tag applied
10 thereto for transit through the multi-stage ATM node of Fig. 1, the example cell being for a 8 bit Utopia device.

Fig. 4D is a diagrammatic view of an example cell having a SPAS tag applied thereto for transit through the multi-stage ATM node of Fig. 1, the example cell being for a 16 bit Utopia device.

Fig. 5 is a schematic view showing a simplification of the multi-stage ATM node
15 of Fig. 1.

Fig. 5A is a schematic view showing demultiplexing points in the multi-stage ATM node as depicted in Fig. 5.

Fig. 5B is a schematic view showing multiplexing points in the multi-stage ATM
20 node as depicted in Fig. 5.

Fig. 5C is a schematic view showing translation points in the multi-stage ATM node as depicted in Fig. 5.

Fig. 5D is a schematic view showing monitoring points in the multi-stage ATM node as depicted in Fig. 5.

Fig. 5E is a schematic view showing activation and deactivation points in the multi-stage ATM node as depicted in Fig. 5.

Fig. 6 is a schematic view showing an origination end and a terminating end of a segment of the multi-stage ATM node of Fig. 1.

5 Fig. 7 is a schematic view showing signaling and cell flow in a performance monitoring example according to a mode of the invention in the multi-stage ATM node of Fig. 1.

Fig. 8 is a diagrammatic view showing an example format of a SPAS performance monitoring control cell according to an embodiment of the invention.

10 Fig. 9 is a diagrammatic view showing transmission of a block of SPAS cells in connection with a performance monitoring operation.

Fig. 10A is a diagrammatic routing depicting an alternative internal routing technique through a multi-stage ATM cell; Fig. 10B is a diagrammatic routing depicting internal routing technique through a multi-stage ATM cell in accordance with a mode of
15 the present invention.

Fig. 11 is a schematic view of an embodiment of the multi-stage ATM node of the present invention having a ring or bus topology.

Fig. 12 is a diagrammatic view of a octet pair employed in a SPAS tag applied for the ringed multi-stage ATM node of Fig. 11.

20 Fig. 13A through Fig. 13F are diagrammatic views showing a SPAS tag according to the invention at different points of travel through the ringed multi-stage ATM node of Fig. 11.

Fig. 14 is a diagrammatic view showing demultiplexing of an ATM cell having AAL2 protocol into an ATM cell having AAL2 prime protocol.

25 Fig. 15 is a diagrammatic view of routing information octet for a traffic cell.

DETAILED DESCRIPTION OF THE DRAWINGS

In the following description, for purposes of explanation and not limitation, specific details are set forth such as particular architectures, interfaces, techniques, etc. in order to provide a thorough understanding of the present invention. However, it will be apparent to those skilled in the art that the present invention may be practiced in other embodiments that depart from these specific details. In other instances, detailed descriptions of well known devices, circuits, and methods are omitted so as not to obscure the description of the present invention with unnecessary detail.

Fig. 1 shows portions of a multi-stage ATM node 20 according to an example embodiment of the invention. In the particular example shown in Fig. 1, multi-stage ATM node 20 is situated in a rack or cabinet of electronic components, the rack comprising plural subracks. For sake of simplification, only five subracks 22 of multi-stage ATM node 20 are shown in Fig. 1, particularly one main subrack 22_M and four access subracks 22_{A1} through 22_{A4}. Space switching components of multi-stage ATM node 20 are hereinafter collectively referred to as "SPAS". The interface between subracks 22 is known as the SILI (SPAS Internal Link Interface) interface 23. Cells which are transmitted through the SPAS are referenced herein as "SPAS cells", and (as hereinafter described) can include both SPAS traffic cells and SPAS performance monitoring control cells (also known as monitoring management cells, SPAS controls cells, or simply control cells).

Each subrack is said to have components comprising an ATM switch mounted thereon. As explained in more detail hereinafter, each subrack 22 comprises a ATM switch core 24. Each switch core 24 has a plurality of switch core ports, particularly a plurality of switch core ingress ports as well as a plurality of switch core egress ports. Each of the switch core ports is connected by an intra-subrack link to a switch port interface module (SPIM) 26. An example of communications between a switch core and switch port interface modules is found in U.S. Patent Application Serial No. 09/188,265, filed November 9, 1998 for "Asynchronous Transfer Mode Switch", which is incorporated herein by reference. Those of the switch port interface modules (SPIM) 26 of access subracks 22_A which are connected to inter-node (e.g., external) links 28 reside on device boards 30.

For illustrative purposes, each of the subracks 22 of multi-stage ATM node 20 is shown as having two switch port interface modules (SPIM) 26_A on an ingress side of switch core 24_A and two switch port interface modules (SPIM) 26_A on an egress side. For example, for subrack 22_{A1} there are provided a first set of switch port interface modules or units (SPIM) 26_{A1-1} and 26_{A1-3} on device boards 30_{A1-1} and 30_{A1-3}, respectively, on ingress side of switch core 24_{A1}. A second set of switch port interface modules or units (SPIM) 26_{A1-2} and 26_{A1-4} are provided on egress side of switch core 24_{A1}. The switch port interface module (SPIM) 26_{A1-4} is connected to another (unillustrated) subrack 22 of multi-stage ATM node 20. The switch port interface module (SPIM) 26_{A1-2} is connected to main subrack 22_M, and particularly to switch port interface module (SPIM) 26_{M-1} thereof. In Fig. 1, for convenience the other access subracks 22_A are shown with comparably referenced components. However, it should be understood that the access subracks 22_A do not need to have identical architectures, but that differing numbers of switch port interface modules (SPIM) 26 can be provided thereon and that such switch port interface modules (SPIM) 26 can be connected differently than shown.

As mentioned above, switch port interface modules (SPIM) 26 serving external links 28 reside, in the illustrated embodiment, on device boards 30. An example device board 30 is illustrated in Fig. 3. As shown in Fig. 3, the device board 30 comprises not only the switch port interface module (SPIM) 26, but one or more SPAS user resources 32. In the particular embodiment shown in Fig. 3, four such SPAS user resources 32 are illustrated. A SPAS user resource 32 can be, for example, a microprocessor, a digital signal processor, ATM or AAL terminating components, or an extension terminal (ET). Extension terminals (ETs) are particularly used for connecting (via an external link 28) multi-stage ATM node 20 to another ATM node in a multi-node network. For functions such as conversion operations and performance monitoring operations hereinafter described, the exchange terminals typically have processors (known as "board processors" or "BPs") mounted thereon. As hereinafter explained, in accordance with the present invention the board processors on device boards 30 at an edge of multi-stage ATM node 20 also function as tagging units, for which reason the board processor/tagging units of Fig. 3 are referenced as BP/TU.

Each of the SPAS user resources 32 connects to the switch port interface module (SPIM) 26 of the device board 30 over a SAI (SPAS Access Interface) interface 34. The switch port interface module (SPIM) 26 includes both hardware and software, and has various components including buffers. Example switch port interface modules (SPIM) 26 are illustrated, for example, in the following United States Patent applications (all of which are incorporated herein by reference): U.S. Patent Application SN 08/893,507 for "Augmentation of ATM Cell With Buffering Data"; U.S. Patent Application SN 08/893,677 for "Buffering of Point-to-Point and/or Point-to-Multipoint ATM Cells"; and U.S. Patent Application SN 08/893,479 for "VP/VC Look-Up Function". The switch port interface module (SPIM) 26 is connected to the switch core 24 of the subrack 22 by ASCI (ATM Switch Core Interface) interface 36.

Fig. 1 primarily shows the subracks 22 of multi-stage ATM node 20. In addition to its subracks 22, the multi-stage ATM node 20 also comprises various managers as shown in Fig. 2. A connection setup manager 50 performs numerous functions, including the forwarding (upon connection set up) of routing tag information to tagging unit BP/TU (see Fig. 3). As hereinafter described, tagging unit BP/TU appends a tag to cells which are to be routed through the ATM node 20, including traffic cells entering the ATM node. As explained subsequently, this tag comprises, e.g., a list of destination addresses for switching of a cell through the ATM node. A node performance monitoring manager 60 controls the monitoring of cell travel through the various switches (e.g., subracks) comprising multi-stage ATM node 20. In addition, a traffic control manager 70 performs various functions, including the function of inserting a quality of service indicator into the tag of the cell.

The connection setup manager 50 and node performance monitoring manager 60 are connected to SPAS by a SPAS Management (SMI) Interface 38 shown in Fig. 2. The SPAS Management (SMI) Interface 38 basically handles the fault, performance, and configuration management of the SPAS. The connection setup manager 50 and node performance monitoring manager 60 can be located at any convenient location in multi-stage ATM node 20, but are preferably located on a main processor MP on a specific device board 30 connected to switch core 24_M (see Fig. 1). The main processor (MP) has the basic function of controlling multi-stage ATM node 20.

Although shown as a single block in Fig. 2, the function of traffic control manager 70 can be performed by various processors located within the SPAS. For example, these functions can be handled by one or more board processors (BPs), e.g., in distributed fashion or by a main processor MP of the multi-stage ATM node 20. Thus, one purpose of the BP can be to handle local traffic control in addition to the local fault and performance monitoring functions.

Before a SPAS cell can enter multi-stage ATM node 20, the traffic control manager 70 must request, over SAI (SPAS Access Interface) interface 34, that a connection be set up between two SAIs with specified quality of service (QoS) and traffic parameters. The connection set up request is received by the SPAS and forwarded to the connection setup manager 50. The connection setup manager 50 responds by providing transfer lists for the connection to two tagging units BP/TUs. The two tagging units BP/TUs which receive transfer lists are those on device boards 30 connected to the two external links 28 involved in the connection. The transfer list is a list of destination addresses to be used for switching and routing of the cell through the SPAS. The transfer list received by one of the two tagging units BP/TUs is used for routing cells in one direction through multi-stage ATM node 20; the transfer list received by the other of the two tagging units BP/TUs is used for routing cells in a reverse direction through multi-stage ATM node 20. The transfer lists are written into conversion tables of the two tagging units BP/TUs.

After connection setup, when an ATM cell participating in the connection is received from an external link 28, tagging unit BP/TU obtains the VPI/VCI and link identifier from the header of the incoming ATM cell. Using these parameters, the tagging unit BP/TU consults its conversion table to obtain the transfer list for the cell. In addition, the traffic control manager 70 adds two parameters, specifically cell size and QoS indicator. The transfer list together with the cell size and QoS indicator are known herein as the "SPAS tag" or "routing tag", or simply "tag". The tagging unit BP/TU adds the SPAS tag to the entire ATM cell to form a SPAS cell, the ATM cell being the SPAS payload (see Fig. 4). The SPAS cell with its SPAS tag 82 is delivered across the SAI (SPAS Access Interface) interface 34 to the SPAS. The SAI (SPAS Access Interface) interface 34 is the user plane interface toward the SPAS.

Thus, before entering at the SAI (SPAS Access Interface) interface 34, the SPAS tag comprising the complete transfer list (together with cell size and QoS indicator) is added to the payload of SPAS cell by tagging unit BP/TU. The SPAS tag is used to propagate the cell through the SPAS. The SPAS tag defines the connection endpoint.

- 5 The SPAS connection can, in its turn, have a number of connections, e.g., ATM connections, multiplexed on it.

An example format of an SPAS cell 80 with a SPAS tag 82 as applied by tagging unit BP/TU is shown in Fig. 4. The SPAS tag 82 precedes the payload 84 of SPAS cell 80. In the illustrated embodiment, the SPAS tag 82 has seven octets. The first octet,
10 known as the service information octet 86, is generated by traffic control manager 70 as mentioned above. The last six octets of SPAS tag 82 are the transfer list 88.

The service information octet 86 has the four following fields: an odd parity field; a cell size field; a quality of service (QoS) field; and a type field. The type field (one bit) has a zero value if the cell is a traffic cell; a one value in the type field
15 indicates that the cell is a control cell or other cell. The QoS field (two bits) has a value of zero if the cell is of the lowest delay priority and a value of three if the cell is of the highest delay priority. The cell size field (four bits) has stored therein a code having a value of from zero to nine. The code is indicative of both total cell size and the size of the cell payload. The meaning of the cell size code is understood with reference to
20 Table 1 as discussed subsequently. Not shown in Table 1 are cell size codes 12, 13, 14, and 15, which are for concatenated cells and which otherwise have the same meanings as codes 0, 1, 2, and 3, respectively.

TABLE 1
CELL SIZE CODING FOR THE SERVICE INFORMATION OCTET

Cell Size Code	0	1	2	3	4	5	6	7	8	9
Total Cell Size	12	18	24	30	36	48	60	60	60	60
SPAS Payload Size	5	11	17	23	29	41	53	53	ATM-Cell	ATM-Cell
AAL2' Payload	2	8	14	20	26	38	45	45	x	x

As indicated above, the transfer list 88 is a list of destination addresses to be used for routing, e.g., switching, of the SPAS cell through multi-stage ATM node 20. In the illustrated embodiment, the transfer list 88 accommodates six octets. As hereinafter explained, the octets of the transfer list 88 are shifted through the transfer list by pop up operations. All of the six octets of transfer list 88 are of identical format, each having a format field; an address field; and a parity bit. The format field (one bit) has a value of zero when the value in the six bit value in the address field is a physical stage address (i.e., the six bit address of the address field corresponds directly to a physical output). The format field (one bit) has a value of one if the value in the address field is to be used for other than a physical address.

When the format field (one bit) of an octet in transfer list 88 has a value of zero, the value in the address field indicates a physical address within multi-stage ATM node 20. In the illustrated examples, these addresses are addresses of switch port interface modules (SPIM) 26, and thus are typically utopia addresses.

When the format field (one bit) of an octet in transfer list 88 has a value of one, the value or logic address (e.g., operation code) in the address field has significance for indicating a certain action. The meanings of these logic addresses depends upon whether the octet is for an even numbered substage or an odd numbered substage. The meanings of these logic addresses for an even numbered stage (e.g., stages 2, 4, and 6) are shown in Table 2. The meanings of these logic addresses for an odd numbered substage (e.g., substages 1, 3, and 5) are shown in Table 3. It is to be noted that the

logic address values in the address field for the format one type octet is only valid at the points where they are processed.

TABLE 2
LOGIC ADDRESS MEANINGS FOR EVEN SUBSTAGES IN TRANSFER
LIST

5

Logic Address Value in Stage Field	Significance
0	best effort broadcast
1	guaranteed broadcast
2	multicast table 1, full multicast table
3	multicast table 2, limited multicast table
4	multicast table 3, limited multicast table
5 - 30	Reserved for limited multicast connections
31	Indicates "null"; the cell shall be terminated if it has reached this far. The code shall be inserted in positions not participating in the routing, e.g., if only one subrack is traversed.
32 - 50	Reserved for multicast tables
51 - 59	Functional address types if the MP/BP has identified itself to the SPIM. If no match, some of the codes could also indicate that the traffic cell shall be broadcasted further.
60-63	Reserved for SPIM HW, indicates that the cell is for SPIM internal use at the termination point. E.g., for fault and performance management (including flow control).

TABLE 3
LOGIC ADDRESS MEANINGS FOR ODD SUBSTAGES IN TRANSFER
LIST

Logic Address Value In Stage Field	Significance
0	best effort broadcast
1	guaranteed broadcast
2	multicast table 1, full multicast table
3	multicast table 2, limited multicast table
4	multicast table 3, limited multicast table
5 - 30	Reserved for limited multicast tables
31	Indicates "null"; the cell shall be terminated if it has reached this far. The code shall be inserted in positions not participating in the routing, e.g., if only one subrack is traversed.
32 - 63	Indicates ring topology with the destination ring subrack addressed..

5

The SPAS tag 82 with its transfer list 88 thus, in the present example, supports a SPAS hierarchic structure of six substages. At each substage up to sixty four outputs can be identified. As indicated above, the transfer list 88 is targeted for (but not limited to) a structure with access subrack switches connected to a main switch. One subrack (e.g., subrack 22) is assumed to consume two of the substages of the transfer list 88. Of an odd and even (e.g, first and second) pair of substages in the transfer list 88, the first substage addresses the output board in the subrack and the second substage addresses the device connected to one or two "multiply utopia" links.

As explained above with reference to the SPAS cell 80 having the format of Fig. 4, each substage or octet in the transfer list 88 has an address field which contains either

15

WO 00/58872 PC1/SE99/02519

a destination, source, or null address. After a destination address in the address field of the octet has been utilized for routing purposes, it is replaced with the address of the address from which the cell came (i.e., the source address). If a connection does not use a full transfer list 88 in order to reach its destination, e.g., only two subracks are traversed, the last two substages are "null". If a switch port interface module (SPIM) 26 detects a source address or null at the top of the transfer list 88, the cell is discarded.

Upon entry into the SPAS, all address fields in transfer list 88 that are to be used for routing through the SPAS are to be populated with valid destination addresses. If the routing chain is shorter than the full possible structure, the address fields of the remaining octets in transfer list 88 are set to null. If a null is at the top of transfer list 88, such is detected and the cell is discarded. As indicated above with reference to the format field of each octet in transfer list 88 (see Fig. 4), the value in the address field can be a physical address or a logical address. The physical address in the address field of an octet of transfer list 88 pinpoints an output and is used for normal point-to-point connections at the substage corresponding to the octet. When the format field of the octet in the transfer list 88 indicates that the value in the address field is a logic address, such logic address is used for various other operations (see Table 2 and Table 3).

At each routing substage the destination address in the address field at the top of the transfer list 88 is used. After the destination address for a particular octet is utilized, the transfer list 88 is pushed up or popped up one step and a source address related to the popped octet is inserted in the address field of the last octet in transfer list 88. That inserted source address is marked so that infinite loops can be detected. If a source address is detected at the top of the transfer list 88, the complete SPAS cell is discarded. The source address signifies the physical source. The format field bit of the popped octet is copied to the last octet in the transfer list 88 (the format field in the last octet thus indicating if the previous value of the address field of the popped octet was a physical or logical address). Odd parity applies for a valid destination address, whereas even parity applies for physical source addresses. If an even parity is detected at a demultiplexing point, it shall be regarded as "null" and the entire SPAS cell is judged invalid.

Fig. 5 is a simplification of the multi-stage ATM node 20 of Fig. 1, showing only main subrack 22_M and the two access subracks 22_{A1} and 22_{A2} . The simplified view of Fig. 5 shows a depiction of a particular route of travel of a block or stream of SPAS cells through multi-stage ATM node 20. The route of travel begins at the ingress SAI (SPAS Access Interface) interface 34 at which the SPAS cell is directed to switch port interface module (SPIM) 26_{A1-1} of access subrack 22_{A1} . From switch port interface module (SPIM) 26_{A1-1} the cell goes through core 24_{A1} to switch port interface module (SPIM) 26_{A1-2} . The switch port interface module (SPIM) 26_{A1-2} applies the cell to a link for transmission to switch port interface module (SPIM) 26_{M-1} of main subrack 22_M . From switch port interface module (SPIM) 26_{M-1} the cell is routed through core 24_M to switch port interface module (SPIM) 26_{M-2} . At switch port interface module (SPIM) 26_{M-2} the cell is applied to a link for transmission to switch port interface module (SPIM) 26_{A2-1} of subrack 22_{A2} . From switch port interface module (SPIM) 26_{A2-1} the cell is routed through core 24_{A2} to switch port interface module (SPIM) 26_{A2-2} . From switch port interface module (SPIM) 26_{A2-2} the cell is applied to the egress SAI (SPAS Access Interface) interface 34, thereby exiting multi-stage ATM node 20. The address fields of the first five of the six octets of the transfer list 88 for the cell routed through multi-stage ATM node 20 in the manner just described would thus have the physical addresses of the following respective SPIMs: 26_{A1-2} ; 26_{M1} ; 26_{M-2} ; 26_{A2-1} ; 26_{A2-2} . The sixth octet carries a valid destination address out on the egress SAI (SPAS Access Interface) interface 34, i.e., out of multi-stage ATM node 20.

Fig. 5A shows demultiplexing points D at which demultiplexing occurs, i.e., at the outlet of each core 24 and at inlets of certain ones of the switch port interface modules (SPIMs) 26 in each subrack 22. The SPAS tag 82 is popped up or pushed up one step at the points labeled "P", e.g., after the demultiplexing points D. In similar manner, Fig. 5B shows the location of multiplexing points ("M") located at the inlet to each core 24 and at the inlet of certain switch port interface modules (SPIMs) 26. The source address is the physical address from the previous multiplexing point closest to the push point. In this way a complete source address list is built up (except for the SAI addresses). The source address transfer list can be use for various purposes, e.g., performance monitoring on any SPAS connection, either end-to-end or any particular segment. The multiplexing points are not controlled by the SPAS tag 82. An SPAS cell is consequently always routed to the next demultiplexing point on which the SPAS tag

82 operates. A SPAS cell across the SAI (SPAS Access Interface) interface 34 must pass at least two multiplexing points M, two demultiplexing points D, and one push point (see Fig. 5A and Fig. 5B). Thus, in the illustrated embodiment, up to five translation points T (e.g., points at which a destination address gets popped up) are possible (see Fig. 5C). As the destination addresses are popped up, the last octet of the transfer list 88 is filled with the source address as described above. The list of source addresses thus being built up in the transfer list 88 can be used to ^{of} quality the SPAS connections that participate in the performance monitoring. *monitor the*

The purpose of the performance monitoring of the invention, as facilitated by node performance monitoring manager 60, is to check the quality of the connection for a block of data, e.g., specified streams of cells routed through at least part of multi-stage ATM node 20. Quality can mean, for example, cell loss and bit error rate. In order to implement the performance monitoring of the present invention, the node performance monitoring manager 60 works in conjunction with various monitoring points established in the SPAS. As explained below, the performance monitoring can be conducted with respect to a cell's entire route of travel through the SPAS, or to one or more segments of the cell's route of travel through the SPAS. The monitoring is made on the payload 84 of the SPAS cell 80 aggregated for all cell sizes in the defined data block and segment.

Fig. 5D shows various potential monitoring points in the SPAS utilized by node performance monitoring manager 60. The potential monitoring points include a monitoring start point MSP; monitoring start/end points MS/EP; and a monitoring end point MEP. Performance monitoring can start at any of the monitoring start point MSP and monitoring start/end points MS/EP inside the SPAS, and can end at any of the monitoring start/end points MS/EP and monitoring end point MEP inside the SPAS. With this in mind, Fig. 5D further illustrates four scenarios of how segments can be defined and monitored in the SPAS by node performance monitoring manager 60. A first such scenario shows all possible (five) SPIM-SPIM paths in the SPAS being defined as a separate segment. A second scenario shows two segments being defined and monitored, each segment being from a first SPIM of a subrack to the first SPIM of a subsequent subrack. The third scenario shows a segment being defined from SPIM 26_{M-1} to SPIM A_{2-2} . The fourth scenario shows a segment being defined as the entire path of the cell through the SPAS, i.e., from SPAS 26_{A1-1} to SPAS 26_{A2-2} .

The shorter a segment, the more SPAS connections can traverse that segment. This is because there is a multiplexing point prior to the monitoring start point and a demultiplexing after the monitoring end point. A maximum segment spanning SAI (ingress) to SAI (egress) can only hold the connections traversing over the SAIs in question.

When a segment is monitored by node performance monitoring manager 60, all SPAS connections and higher layer connections (e.g., ATM connections) that are multiplexed on that segment are monitored so long as they have the specified quality of service (QoS), a valid physical destination address at the start point, and a valid source address at the end point for the substages of the transfer list 88 which comprise the segment. To be valid, a physical address must comprise at least one substage.

In the monitoring operation performed under control of node performance monitoring manager 60, a block of data is bounded by a not participating start cell and a stop cell, which are described in more detail subsequently in connection with SPAS control cells. The quality on all valid cells (traffic cells and otherwise) between the start cell and the stop cell is monitored.

The establishment of a monitoring operation supervised by node performance monitoring manager 60 has three phases – an activation phase; a monitoring phase; and a reporting phase. Prior to discussing each phase, mention is first made regarding activation and deactivation points for the performance monitoring. Fig. 5E specifically shows potential activation and deactivation points in connection with the example embodiment previously discussed. Fig. 5E shows location of the following potential points: an activation point (AP); activation/reporting points (A/RP); and a reporting point (RP). From Fig. 5E it can be seen that SPIM 26_{A1-1} can serve only as an activation point (AP); that SPIM 26_{A2-2} can serve only as a reporting point (RP); and that the SPIMs 26_{A1-2} and 26_{M-2} can serve either as an activation point or a reporting point, e.g., can serve as activation/reporting points (A/RP). An activation point is a point that can, under supervision of node performance monitoring manager 60, initiate performance monitoring and define the segment size. An activation point can only define segments starting in the same SPIM. Similarly, a reporting point can only operate on a monitoring end point in the same SPIM.

The function of activation and reporting points is understood with reference to Fig. 6. Fig. 6 particularly shows two SPIMs 26 which are situated at endpoints of a monitored segment. In Fig. 6, SPIM 26₀ is referred to as the originating SPIM, while SPIM 26_T is referenced as the terminating SPIM. The segment monitored could be any possible segment in the SPAS, e.g., any of the segments shown in Fig. 5D, for example. Each of the SPIMs 26₀ and 26_T is shown as having a board processor (BP) and hardware (HW). The SPIM 26₀ is shown as having an activation point AP in its board processor (BP) and a monitoring start point (MSP) in its hardware. Similarly, SPIM 26_T has a reporting point RP in its board processor (BP) and a monitoring end point (MEP) in its hardware. The SPIMs 26₀ and 26_T are shown as residing on their respective subracks 22, each subrack 22 having a switch core (ASCM) 24.

Since the transfer list 88 is built up as a pair of destination addresses, the monitoring control signals must traverse the ASCM (switch core) in both subracks in order to span the SILI (SPAS Internal Link Interface) interface 23 [see Fig. 6]. The board processors (BPs) in SPIM 26₀ and SPIM 26_T conduct the monitoring, under supervision of node performance monitoring manager 60. As shown in Fig. 6, all control signals in phase 1 (the activation phase) and in phase 3 (the reporting phase) are carried directly between the board processors (BPs) of SPIM 26₀ and SPIM 26_T. In phase 2 (the monitoring phase), the start and stop signals (described in more detail subsequently with as "start" and "stop" SPAS control cells) are recognized by the monitoring start point MSP and the monitoring end point MEP as well. The monitored connection is only recognized by the monitoring points, and only for the duration of the monitoring.

The board processors of originating SPIM 26₀ and terminating SPIM 26_T communicate with one another using special cells known as SPAS performance monitoring control cells, also known as "monitoring management cells", "SPAS control cells", or simply "control cells". An example format of a SPAS performance monitoring control cell is shown in Fig. 8. In the illustrated embodiment, each SPAS performance monitoring control cell is thirty octets in length. In such embodiment, the SPAS performance monitoring control cell has five fields: a header field 8-1; a performance monitoring code field 8-2; a correlation field 8-3; a data field 8-4; and, a CRC field 8-5. All but the header field 8-1 and the data field 8-4 are but one octet in

length, the header field 8-1 being seven octets and the data field 8-4 being twenty octets. The header field 8-1 of the SPAS performance monitoring control cell is the very tag applied by the tagging unit (BP/TU) which also tags traffic cells.

As indicated by the value in its performance monitoring code field 8-2, a SPAS performance monitoring control cell can be one of three types. If the performance monitoring code has a value of zero, the SPAS performance monitoring control cell indicates that performance monitoring has been activated. An activation SPAS performance monitoring control cell carries, in its data field 8-4, comparison data that shall be used by the monitoring end point and indicates which part(s) of the transfer list 88 is being monitored and the quality of service (QoS) parameter for monitoring. By “comparison data” is meant contents of an entire transfer field that is to be used by terminating SPIM 26_T for setting up its monitoring activity.

If the performance monitoring code has a value of one, the SPAS performance monitoring control cell is a “respond” cell sent from terminating SPIM 26_T and indicating whether terminating SPIM 26_T accepts the monitoring function or not. If the performance monitoring code has a value of one, the SPAS performance monitoring control cell is a “result” cell sent from terminating SPIM 26_T to originating SPIM 26_O and having monitored data collected at the monitoring end point (MEP) in its data field 8-4.

Fig. 7 provides an performance monitoring example which illustrates all three phases – the activation phase; the monitoring phase; and the reporting phase. As event 7-1, node performance monitoring manager 60 sends an initiate performance monitoring signal to the board processor of originating SPIM 26_O. Event 7-1 starts the activation phase.

As part of the activation phase, the board processor (BP) of originating SPIM 26_O sends an activation SPAS performance monitoring control cell as event 7-2 to the board processor (BP) of the terminating SPIM 26_T. The activation SPAS performance monitoring control cell has a performance monitoring code of zero in its field 8-2 (see Fig. 8), and carries in its data field 8-4 an entire transfer list to be used as comparison data for the monitoring, as well as the quality of service (QoS) parameter involved in the monitoring. At the same time, originating SPIM 26_O prepares the hardware

resources at the monitoring start point (MSP) for the monitoring phase (as indicated by event 7-3). The preparation of event 7-3 means that the monitoring start point (MSP) begins (1) to look for start control cells and stop control cells which respectively start and stop monitoring on the segment to be monitored, and (2) [between the start and stop control cells] to look for cells which have, in their transfer list 88, both the specific transfer list pattern specified as the comparison data upon activation and the specified QoS parameter.

Upon receipt of the activation SPAS performance monitoring control cell sent as event 7-2, the board processor (BP) of terminating SPIM 26_T analyzes the content of the activation SPAS performance monitoring control cell (particularly data field 8-4 which includes the entire transfer list and quality of service indicator [see Fig. 8]) and determines whether terminating SPIM 26_T can participate in the performance monitoring requested by originating SPIM 26_O. Reasons for not being able to participate in the performance monitoring may be lack of resources at the targeted terminating SPIM 26_T, or that the existing resources at the switch port interface module (SPIM) 26 are occupied either by other active performance monitoring or other activities. If terminating SPIM 26_T determines that it can participate in the performance monitoring, switch port interface module (SPIM) 26 prepares its resources for such participation, as indicated by event 7-3. The preparation of event 7-3 involves the board processor (BP) of terminating SPIM 26_T advising the hardware (HW) of terminating SPIM 26_T that the monitoring end point (MEP) thereof should look for SPAS tags 82 with a certain source address and QoS in the selected part of the transfer list 88 for participating cells, as well as to be on the lookout for the specific code in the transfer list 88 indicating to start and stop the monitoring. The terminating SPIM 26_T also sends a response SPAS performance monitoring control cell to originating SPIM 26_O as event 7-4, indicating whether or not the terminating SPIM 26_T can participate in the performance monitoring. As indicated previously, the response SPAS performance monitoring control cell has a value of one in its performance monitoring code field 8-2 (see Fig. 8).

Upon receipt of the response SPAS performance monitoring cell from terminating SPIM 26_T the monitoring phase begins (see Fig. 7). In the monitoring phase generally, the monitoring start point (MSP) generates check data for a block of

SPAS cells which have a common physical destination address in the transfer list 88 all the way to the specified monitoring end point (MEP), e.g., the terminating SPIM 26_T . In the monitoring phase, both the monitoring start point (MSP) at originating SPIM 26_0 and the monitoring end point (MEP) at terminating SPIM 26_T are prepared to look in
5 the SPAS tag 82 of received SPAS cells, and particularly in transfer list 88 thereof, for their specific pattern. The pattern sought by the monitoring start point (MSP) is not the same as the pattern sought by the monitoring end point (MEP), since the monitoring start point (MSP) looks for a particular destination address in the appropriate octet of transfer list 88 while the monitoring end point (MEP) looks for a source address that
10 identifies originating SPIM 26_0 .

Describing now the monitoring phase in more detail, under supervision of node performance monitoring manager 60, the originating SPIM 26_0 issues a SPAS control cell that includes two instances of a start code, i.e., a "62" value in the address field of two octets in transfer list 88 of SPAS tag 82 (see Table 2). In which two of the octets
15 the start code is included depends upon and corresponds to the location of the monitoring stop point. In other words, in the transfer list, the start code is substituted in two positions (e.g., two substages) corresponding to the SPIMs which are supposed to act on the start code. When the SPAS control cell containing the start code is detected at the monitoring start point (MSP), the monitoring begins at originating SPIM 26_0 .
20 Moreover, having discovered from the start code that the monitoring is to begin, the monitoring start point (MSP) then removes the first instance of the start code from the transfer list and substitutes therefor the address of the monitoring start point. In connection with this substitution, the SPIM looks at a correlated value that is stored during activation by its local board processor (see Fig. 6) associated with this
25 correlation value is the real tag value for the position in the SPIM. In this way, the address of the monitoring start point (MSP) continues with the SPAS cell (which still includes a second instance of the start code).

Fig. 7 shows as event 7-5 the transmission of this same SPAS control cell with the start code from originating SPIM 26_0 to terminating SPIM 26_T . When that same
30 SPAS control cell containing the second instance of the start code is detected at the monitoring end point (MEP), the monitoring begins at terminating SPIM 26_T . The presence of the start code (second instance of) in a position of the transfer list

corresponding to the monitoring end point (MEP) causes the terminating SPIM to realize that it is the monitoring end point. In like manner as was done at the monitoring start point (MSP), the terminating SPIM replaces the instance of the start code with the address of the terminating SPIM so that the SPAS cell can now continue with a complete transfer list.

The flow of further SPAS cells (which can be traffic cells and can include other SPAS control cells) from originating SPIM 26_0 to terminating SPIM 26_T is indicated by event 7-6 (see also Fig. 9). The SPAS cells can be any of the approved sizes (see the description of the cell size field in service information octet 86 of Fig. 4). Monitoring data is generated both at originating SPIM 26_0 and terminating SPIM 26_T for each SPAS cell with the comparison data pattern in its transfer list 88.

While the SPAS cells are transmitted from originating SPIM 26_0 to terminating SPIM 26_T as part of event 7-6 of the monitoring operation (see Fig. 7 and Fig. 9), monitoring data is maintained at both originating SPIM 26_0 and terminating SPIM 26_T . The monitoring data can be of several conventional forms, but preferably is cell count and/or integrity check of the content of the payload 84 by means of a total checksum for all such cells.

Under supervision of node performance monitoring manager 60, at an appropriate time the originating SPIM 26_0 issues a SPAS control cell that includes a stop code, i.e., a "63" value in the address field of two octets of transfer list 88 of SPAS tag 82 (see Table 2). As with the two instances of the start code, the two instances of the stop code occur in substages of the transfer list that correspond to the monitoring start point (MSP) and the monitoring end point (MEP). When a SPAS control cell bearing the stop code is detected at monitoring start point (MSP), originating SPIM 26_0 stops its gathering of monitoring data, and replaces the first instance of the stop code with the monitoring start point (MSP) address. The SPAS control cell still bearing the second instance of the stop code is transmitted on to terminating SPIM 26_T , as indicated by event 7-7 in Fig. 7. When the SPAS control cell bearing the second instance of the stop code is received at monitoring end point (MEP), terminating SPIM 26_T also ceases its gathering of monitoring data and replaces the second instance of the stop code with the monitoring end point (MEP) address. In essence, the performance monitoring data collected at both originating SPIM 26_0 and terminating SPIM 26_T is frozen. The board

processor (BP) of originating SPIM 26₀ and the board processor (BP) of terminating SPIM 26_T both generate monitoring data results, as shown by event 7-8 and event 7-9, respectively. In generating the monitoring data results, the board processors (BPs) of the SPIMs read registers that have the monitoring data stored therein.

5 Upon completion of its generation of monitoring data results, the board processor (BP) of terminating SPIM 26_T issues a results report, also known as the termination unit report, to the board processor (BP) of originating SPIM 26₀ as event 7-10. The results report is included in the report SPAS performance monitoring control cell as previously described. In particular, the data field 8-4 of the SPAS performance monitoring control
10 cell includes the monitoring data results collected at the monitoring end point (MEP).

 Upon receipt of the report SPAS performance monitoring control cell from terminating SPIM 26_T, the board processor (BP) of originating SPIM 26₀ performs an analysis to compare the received results from terminating SPIM 26_T with its own results, as indicated by event 7-11. The originating SPIM 26₀ develops a conclusion
15 based on its analysis, and sends a conclusion report to node performance monitoring manager 60, as indicated by event 7-12. As indicated previously, node performance monitoring manager 60 may be located in a main processor of multi-stage ATM node
20.

 As an alternative to the foregoing reporting scheme, both terminating SPIM 26_T
20 and originating SPIM 26₀ can separately forward their performance monitoring data results to node performance monitoring manager 60, so that node performance monitoring manager 60 can perform the analysis.

 Returning now to Table 1, each cell size code of the service information octet 86 specifies a set of sizes, including a total cell size, a SPAS payload size, and an AAL2'
25 payload size. AAL2' (also written AAL2 prime) is a special protocol which is described in United States Patent Application Serial No. 09/188,102, filed November 9, 1998, entitled "Asynchronous Mode Transfer System", which is incorporated herein by reference. AAL2 prime (AAL2') requires that AAL2 packets carried in the ATM cell payload be whole packets and that the ATM payload not have an AAL2-type start field.
30 Preferably, in the AAL2 prime protocol only one whole AAL2 packet is carried per ATM cell payload. It will be recalled that AAL2 is a standard defined by ITU

recommendation I.363.2. An AAL2 packet comprises a three octet packet header, as well as a packet payload. The AAL2 packet header includes an eight bit channel identifier (CID), a six bit length indicator (LI), a five bit User-to-User indicator (UUI), and five bits of header error control (HEC). The AAL2 packet payload, which carries user data, can vary from one to forty-five octets. Fig. 14 is a diagrammatic view showing demultiplexing of an ATM cell having AAL2 protocol into an ATM cell having AAL2 prime protocol.

As reflected by Table 1, cell size codes 1 - 6 in a service information octet 86 are used for AAL2' format (or perhaps another ATM format, if desired). Fig. 4A shows the format of a SPAS cell 80A which carries an AAL2' packet. As with all other cells transmitted through multi-stage ATM node 20, SPAS cell 80A of Fig. 4A has a SPAS tag 82. The SPAS tag 82 has the same seven octet format as shown in Fig. 4. The SPAS tag 82 is followed by the AAL2' packet 400A, which includes the AAL2' header 402A and an AAL2' packet payload 404A. The AAL2' packet payload 404A can carry up to 45 octets. If an AAL2 packet multiplexed into the AAL2' protocol requires greater than 45 octets, the AAL2 packet must be segmented into two AAL2' packets. The first packet uses a LI code (see Fig. 4A) above 45 (e.g., 48) to indicate a fixed predefined AAL2' packet size, e.g., 32 octets. The LI code of the last AAL2' packet indicates the real size of the last of the two packets. Upon receipt of the two AAL2' packets at a receiving side, they are reassembled into a unit. The AAL2' header is protected by an odd parity bit.

Cell size code 7 in a service information octet 86 (see Fig. 4 and Table 1) is used for yet another protocol, AAL2'' (also written AAL2 double prime). Fig. 4B illustrates a SPAS cell 82B which invokes the AAL2'' protocol, and also illustrates a AAL2'' packet 400B having the AAL2'' protocol. In the AAL2'' protocol, an AAL2' packet (such as AAL2' packet 400A) is carried in an ATM cell, and the ATM-VCI is used to indicate the connection. SPAS cell 80B begins with the SPAS tag 82 (same format as Fig. 4), and is followed by the AAL2'' packet 400B. The AAL2'' packet 400B includes the ATM header (5 octets, including 12 bits of the active ATM-VCI) and the AAL2' packet 400A. In the AAL2'' protocol, the AAL2' VCI is copied to the twelve least significant bits of the ATM VCI. The most significant bits are set to zero, as are the VPI, PTI, and CLP.

The AAL2'' protocol allows an easy protocol transformation between AAL2' and AAL2'' that can be accomplished in hardware the switch port interface module (SPIM) 26. This transformation can be useful in those cases where the external equipment/component only recognizes sixty octet cells with an ATM header and not AAL2'. Examples of those cases are an AAL5-SAR (Segmentation and Reassembly) component interfacing a main processor (MP) or an Exchange Terminal (ET) [in the later case if for some reason the AAL2'' is used on the external ATM links instead of AAL2].

For cell size codes 7 - 9, an additional adaptation must be made to the actual Utopia implementation, depending on whether the Utopia device is eight bits or sixteen bits wide. Fig. 4C shows an example ATM cell format for a 8 bit Utopia, which includes the SPAS tag 82 and has a total SPAS cell size of 60 octets. The multi-stage ATM node 20 transfers the entire ATM cell transparently between the two endpoints. Fig. 4D, on the other hand, shows an example ATM cell format for a 16 bit Utopia, which includes the SPAS tag 82 and has a total SPAS cell size (over the SAI interface) of 62 octets. For the Fig. 4D cell, octet 8 and octet 14 are removed by the multi-stage ATM node 20 during internal transport. Octets 9 - 13 and 15 - 62 are carried transparently (since the multi-stage ATM node 20 will convert between the two Utopia formats when needed).

Cell size code 8 (see Fig. 4 and Table 1) is used to denote a transparent ATM cell. Cell size code 9 used to denote an ATM AAL5 cell that could be subject to Early Packet Discard (EPD). Cell size codes 12 - 15 are for internal usage of multi-stage ATM node 20, while cell size code 10 is reserved for future use.

The multi-stage ATM node 20 can also be structured to have a bus or ring configuration, as illustrated in Fig. 11 for example. The ringed multi-stage ATM node 20R of Fig. 11 comprises n number of subracks, particularly subracks 22_{R0} through 22_{Rn}. The subracks 22_R are connected by a bus or ring R. As in previous embodiments, Each subrack 22_R has a switch core 24 connected intermediate two sets of switch port interface modules (SPIM) 26. For example, subrack 22 has switch core 24_{R0}; switch port interface module (SPIM) 26_{R0-1} (also labeled "SPIM #2"); and switch port interface module (SPIM) 26_{R0-2} (also labeled "SPIM #0") at address (adr) 0. At address adr =1 the SPIM 26_{R2} is connected to ring R. The switch port interface module

(SPIM) 26_{R0-1} is shown (in manner similar to Fig. 1) as being situated on a device board 30_{R0-1} . For sake of simplification, in the ensuing discussion, reference to the various switch port interface modules (SPIMs) 26 will simply be to SPIM #0 and SPIM #2 on subrack 22_{R0-1} , to SPIM 26_{R1-2} on subrack 22_{R1-2} , and to SPIM #5 and SPIM #28 on subrack 22_{R2} , as shown in Fig. 11.

When a bus or ring topology is employed for a multi-staged ATM node, the SPAS tag 82 has two adjacent octets combined to yield an octet pair which represents the bus or ring. Such an octet pair is shown in Fig. 12 as octet pair 1200. The following fields are provided in the first octet of octet pair 1200: format field 1202; type field 1204; (bus or ring) destination address field 1206; and odd parity field 1208. The type field 1204, when set to "one" indicates that a bus or ring topology is invoked. A bus or ring topology can support, in the illustrated embodiment, up to 32 subracks. The destination address field 1206 contains an address of a target ring subrack. The following fields are provided in the second octet of octet pair 1200: "L" field 1210; "SEQ" field 1212; (bus or ring) source address field 1214; and even parity field 1216. For the "L" field 1210, a "1" indicates a logic address (in which case the logic address occupies the destination giving 2^{10} logic combinations of destination and source address including broadcast, multicast, and resource shift). The "SEQ" field 1212 can be used as a sequence counter on a link set or used to extend the logic address field, if necessary.

In an example scenario of a ring-employing multi-stage ATM node $20R$, typically an ingress subrack is attached to the ring and an egress subrack is attached to the ring. Such example scenario is shown in Fig. 11, wherein subrack 22_{R0} serves as the ingress subrack and subrack 22_{R2} serves as the egress subrack, both subrack 22_{R0} and subrack 22_{R2} being connected to ring R. Each subrack is connected to ring R over a SILI (SPAS Internal Link Interface) interface $23R$. The SILI $23R$ is connected as a closed ring connecting from subrack to subrack. The physical wires of ring R are bidirectional in the illustrated embodiment, but it should be understood that unidirectional wires can also be employed.

In the example scenario illustrated in Fig. 11, and also depicted with reference to Fig 13A - Fig. 13F, a SPAS cell is to be sent from SPIM #2 (i.e., SPIM 26_{R0-1}) on subrack 22_{R0} to address $adr=4$ (which is connected to SPIM #5 on subrack 22_{R2}). Fig.

11 shows six specific points of travel, particularly points A - F for such SPAS cell. The SPAS tag 82 for travel point A is shown in Fig. 13A; the SPAS tag 82 for travel point B is shown in Fig. 13B; and so forth. Therefore, Fig. 13 A - Fig. 13F are referenced in the ensuing detailed description of the travel of the cell from point A (SPIM #2) to point F (the device having $adr=4$). Since the octets of SPAS tag 82 are popped in revolving fashion, the first substage octet is referenced as octet O1, the second substage octet as O2, and so forth.

At point A in Fig. 11, the SPAS tag 82 is as shown in Fig. 13A. Thus, SPAS tag 82 has its middle two octets (O3 and O4) of the transfer list in the form of the octet pair 1200 mentioned above (see Fig. 12). After the cell leaves point A, core 24_{R0} routes the cell to SPIM #0 of subrack 22_{R0} at address $adr=0$. In so doing, core 24_{R0} substitutes the source address (the address of SPIM #2) for the destination address in the top octet (octet O1) of transfer list 88.

When the cell reaches ingress to SPIM #0, the SPIM #0 changes the parity of the top octet (octet O1) of transfer list 88 from odd to even parity and then pops the top octet (octet O1) from the top of transfer list 88 to the bottom of transfer list 88. Thus, at point B (inside SPIM #0), the SPAS tag 82 appears as in Fig. 13B. The SPIM #0 looks at the then uppermost octet (octet O2) in the transfer list 88, and determines therefrom that the next physical destination address is $adr=1$. The physical destination address $adr=1$ is the address of ring R. It should be kept in mind that, although unillustrated, several other rings could be connected to SPIM #0 (or any other SPIM, for that matter). SPIM #0 substitutes its address in the top octet (octet O2) of transfer list 88 for the destination address, changes the parity, and then pops the top octet (octet O2) of transfer list 88, so that octet O2 is goes to the bottom of transfer list 88 as shown in Fig. 13C.

The cell with SPAS tag 82 as shown in Fig. 13C travels on ring R until it is accepted by the subrack identified by the destination address field in the top of the transfer list 88. The travel on ring R is represented as point C in Fig. 11. In subrack 22_{R1} the cell is transferred transparently through its core 24_{R1} , since the format and type code in octet pair 1200 indicate that core 24_{R1} is not to touch the cell.

When the cell reaches its ring destination address at SPIM #28, which is at ring address 2 (being subrack 22_{R2}), the cell is accepted by SPIM #28. Moreover, SPIM #28

changes the parity from odd to even of the octet pair 1200 at the top of transfer list 88, e.g., octets O3 and O4, and pops the octet pair 1200 from the top to the bottom of transfer list 88. Therefore, at point D as the cell enters core 24_{R2} from SPIM #28, the SPAS tag 82 appears as in Fig. 13D.

Core 24_{R2} exchanges the contents of the ring destination address and the ring source address fields 1206 and 1214 of the octet pair 1200 at the bottom of transfer list 88, e.g., octets O3 and O4. After the exchange, at point E the SPAS tag 82 has the appearance shown in Fig. 13E. Then core 24_{R2} delivers the cell to the address of the target SPIM #5, e.g., adr=5, which now appears in the uppermost octet (octet O5) of transfer list 88.

The target SPIM #5 puts the uppermost octet (octet O5) at the bottom of the transfer list 88, and changes the parity from odd to even. The target SPIM #5 then transfers the cell, with its SPAS tag 82 at point F appearing as in Fig. 13F, to the addressed application (e.g., the device having adr=4).

Thus, the foregoing describes the routing of a cell through a multi-stage ATM node 20 having a ring topology, and shows the use of octet pair 1200 as well as the popping of octets in transfer list 88 (which is common to all topologies). Thus, the SPAS tag 82 of the present invention accommodates a multi-stage ATM node 20 having a ring or bus topology.

The SPAS tag 82 of the present invention also allows for multicasting or broadcasting of traffic cells. Whether a cell is a traffic cell is determined by the type field of the service information octet 86 (see Fig. 4 and the description thereof). For a traffic cell, the format field of each octet of payload 84 is interpreted as being a cast field (see Fig. 15). The cast field indicates whether (1) the cell is unicast, e.g., the destination address is binary encoded [when the value in cast field is 0], or (2) whether the cell is to be broadcast or multicast. If the cast field indicates multicast or broadcast, the destination address is logical and is interpreted in accordance with Table 4.

TABLE 4

DESTINATION FIELD INTERPRETATION IN TRANSFER LIST FOR
TRAFFIC CELLS

Destination Field Value	Significance
0	broadcast on not occupied crosspoints
1	broadcast, all crosspoints loaded regardless of previous state
2	multicast table 1, full multicast table
3	multicast table 2, limited multicast table
4	multicast table 3, limited multicast table
5 - 30	Reserved for limited multicast connections
31	Indicates "null"; the cell shall be terminated if it has reached this far.
32 - 63	Reserved for ring topology (the core routes the cell to the address stated in the predefined register or to the source, the routing information octet is not changed)

While the present description has utilized an example of multi-stage ATM node 20 having as many as six stages through which an SPAS cell can travel, and consequently a transfer list 88 of six octets, it should be understood that the principles of the invention are not limited to this particular example. Rather, a greater or lesser
5 number of stages may be present in multi-stage ATM node 20, and consequently the length of transfer list 88 is also variable.

Similarly, other parameters described herein are not critical but instead can assume other values in other embodiments. For example, the size of a SPAS performance monitoring control cell as being thirty octets is just one illustrative but
10 variable parameter. In the same vein, if the SILI interface 23 and the SAI interface are present at the same SPIM, the address field of an octet of transfer list 88 can be divided.

In the illustrations provided herewith, e.g., Fig. 5 and Fig. 5A - Fig. 5D, cell travel has been depicted as from right to left. The reader will appreciate, however, that cells are also traveling from left to right, and that such cells are also routed and
15 monitored in accordance with the principles of the invention.

The reader will also understand that times are required, e.g., at the various SPIMs, in order to make the performance monitoring procedure robust, e.g., for the signals requiring a response a local timer is set. The provision and use of timers is well within the ken of the person skilled in the art.

20 The present invention provides an advantageous method for performance monitoring of a multi-staged ATM node. Any connection through the node, or segments thereof, can be monitored in order to detect, e.g., early degradation in performance.

The present invention is both scalable and upgradable. Moreover, the
25 performance monitoring capabilities described herein can be introduced in part of gradually, if desired. For example, the hardware employed can be designed at low cost if desired to handle only one active segment start or end point at a time. Subsequent versions can be used to handle a multitude of segments simultaneously.

The advantages of using the routing tag (e.g., SPAS tag 82) of the invention are manifold. In particular, no VPI/VCI conversion is performed for routing the ATM cell through plural stages of multi-stage ATM node 20 subsequent to the appending of the routing tag. The number of internal control paths is also reduced, and connection set up
5 time is shortened. Such is attributable, e.g., to the fact that mapping between VPI/VCI values is not necessary on internal links within multi-stage ATM node 20. The superiority of the internal routing of the present invention, as represented by Fig. 10B, is understood, e.g., with reference to an alternative technique shown in Fig. 10A. In Fig. 10A, for sake of simplicity extension terminals (ETs) are shown as connected to the
10 switch cores 24, rather than the entire device boards 30 with switch port interface modules (SPIMs) 26. Fig. 10A shows an alternative technique in which VPI/VCI conversion is performed between each of three subracks, particular subrack A, subrack B, and subrack C.

The routing tag of the invention can be utilized in other implementations, and
15 can be used as a token and passed on to the next node in a chain or in a ring if no match is detected. In such case the tag that is passed on may be extended with a node address. In such case, extension terminals (ETs) must be setup in advance in order to know where to pass the cell if no match is detected.

While the invention has been described in connection with what is presently
20 considered to be the most practical and preferred embodiment, it is to be understood that the invention is not to be limited to the disclosed embodiment, but on the contrary, is intended to cover various modifications and equivalent arrangements included within the spirit and scope of the appended claims.

WHAT IS CLAIMED IS:

1. A method of routing an ATM cell through a multi-stage ATM node, each stage of the multi-stage ATM node comprising an ATM switch, the method comprising, appending to at least a payload of a received ATM cell a routing tag, the routing tag comprising routing information for routing the payload of the received ATM cell
5 through plural stages of the multi-stage node whereby no VPI/VCI conversion is performed for routing the ATM cell through the plural stages of the multi-stage node subsequent to the appending of the tag.
2. The method of claim 1, wherein the routing tag comprises routing information for routing the payload of the received ATM cell entirely through the multi-stage node whereby no VPI/VCI conversion is performed for routing the ATM cell through the multi-stage node subsequent to the appending of the tag.
3. The method of claim 1, wherein the routing information comprises a list of destination addresses.
4. The method of claim 3, wherein each stage of the multi-stage node comprises a switch core connected to a first set of interface units and a second set of interface units, and wherein the method comprises including in the routing tag, as the destination addresses, physical addresses for one of the first set of interface units and one of the
5 second set of interface units for each stage of the multi-stage node.
5. The method of claim 3, wherein the method further comprises exchanging one of the destination addresses in the list with a source address when the cell is in route to or received at a destination address specified by the list, the source address being an address from which the cell was routed to the destination address.
5
6. The method of claim 3, wherein a parity value is associated with each destination address in the list of destination addresses comprising the list, and wherein the method further comprises changing the parity value from a first value to a second value when the cell is in route to or received at a destination address specified by the
5 list.

7. The method of claim 3, wherein the node has plural substages, wherein the routing information comprises list with an entry for each substage, and wherein the method further comprises changing the relative order of the entry for a substage when the cell is in route to or received at a destination address associated with one of the plural substages.

8. The method of claim 3, wherein the node has plural substages, wherein the routing information comprises list with an entry for each substage, the list including a parity value for each entry in the list, and wherein the method further comprises, when the cell is in route to or received at a destination address specified by a selected entry in the list:

exchanging the destination address in the entry with a source address from which the cell was routed to the destination address;

changing the parity value in the entry from a first value to a second value;

changing the relative order of the entry in the list.

9. The method of claim 8, wherein the step of changing the relative order of the entry in the list comprises moving the entry from a top of the list to a bottom of the list.

10. The method of claim 1, wherein the addresses are utopia addresses.

11. The method of claim 1, wherein the multi-stage node comprises three stages, and wherein the step of appending the routing tag comprises appending six addresses of the multi-stage node.

12. The method of claim 1, wherein at least some of the addresses are utopia addresses.

13. The method of claim 1, wherein the step of appending the routing tag comprises appending the routing tag upon ingress of the ATM cell to the multi-stage node.

14. The method of claim 1, further comprising, upon connection set up, sending information to be used for the routing tag from a main processor of the multi-stage node to a processor which performs the appending step.

15. The method of claim 1, wherein at least some of the plural stages of the multi-stage ATM node are connected by a bus or a ring, and wherein the appending of the routing tag comprises appending two entries in the tag, the two entries identifying a ring destination address and a ring source address.

16. A multi-stage ATM node comprising:

multiple ATM switches connected together, each ATM switch forming a stage and including an ATM switch core connected between a first set of interface units and a second set of interface units;

5 a tagging unit which appends to at least a payload of a received ATM cell a routing tag, the routing tag comprising routing information for routing the payload of the received ATM cell through plural ones of the ATM switches of the multi-stage node whereby no VPI/VCI conversion is performed for routing the ATM cell through the plural ones of the ATM switches of the multi-stage node subsequent to the appending of
10 the tag.

17. The apparatus of claim 16, wherein the routing tag comprises routing information for routing the payload of the received ATM cell entirely through the multi-stage node whereby no VPI/VCI conversion is performed for routing the ATM cell through the multi-stage node subsequent to the appending of the tag.

18. The apparatus of claim 16, wherein the routing information comprises a list of destination addresses.

19. The apparatus of claim 18, wherein the routing tag includes, as the destination addresses, physical addresses for one of the first set of interface units and one of the second set of interface units for each stage of the multi-stage node.

20. The apparatus of claim 19, wherein at least some of the addresses are utopia addresses.

21. The apparatus of claim 16, wherein the multi-stage node comprises three stages, and the routing tag comprises six addresses of in the multi-stage node.
22. The apparatus of claim 21, wherein at least some of the addresses are utopia addresses.
23. The apparatus of claim 22, wherein the addresses are addresses of route-related ones of the interface units.
24. The apparatus of claim 16, wherein the tagging unit is located at an edge of the multi-stage node and appends the routing tag upon ingress of the ATM cell to the multi-stage node.
25. The apparatus of claim 16, further comprising a connection set up processor which, upon connection set up, sends to the tagging unit information to be used for the routing tag.
26. The apparatus of claim 25, wherein the connection set up processor is located at a stage of the multi-stage node which serves as a main switch.

THIS PAGE BLANK (UBPTO)





Fig. 1

THIS PAGE BLANK (USPTO)

Fig. 2



THIS PAGE BLANK (USPTO)


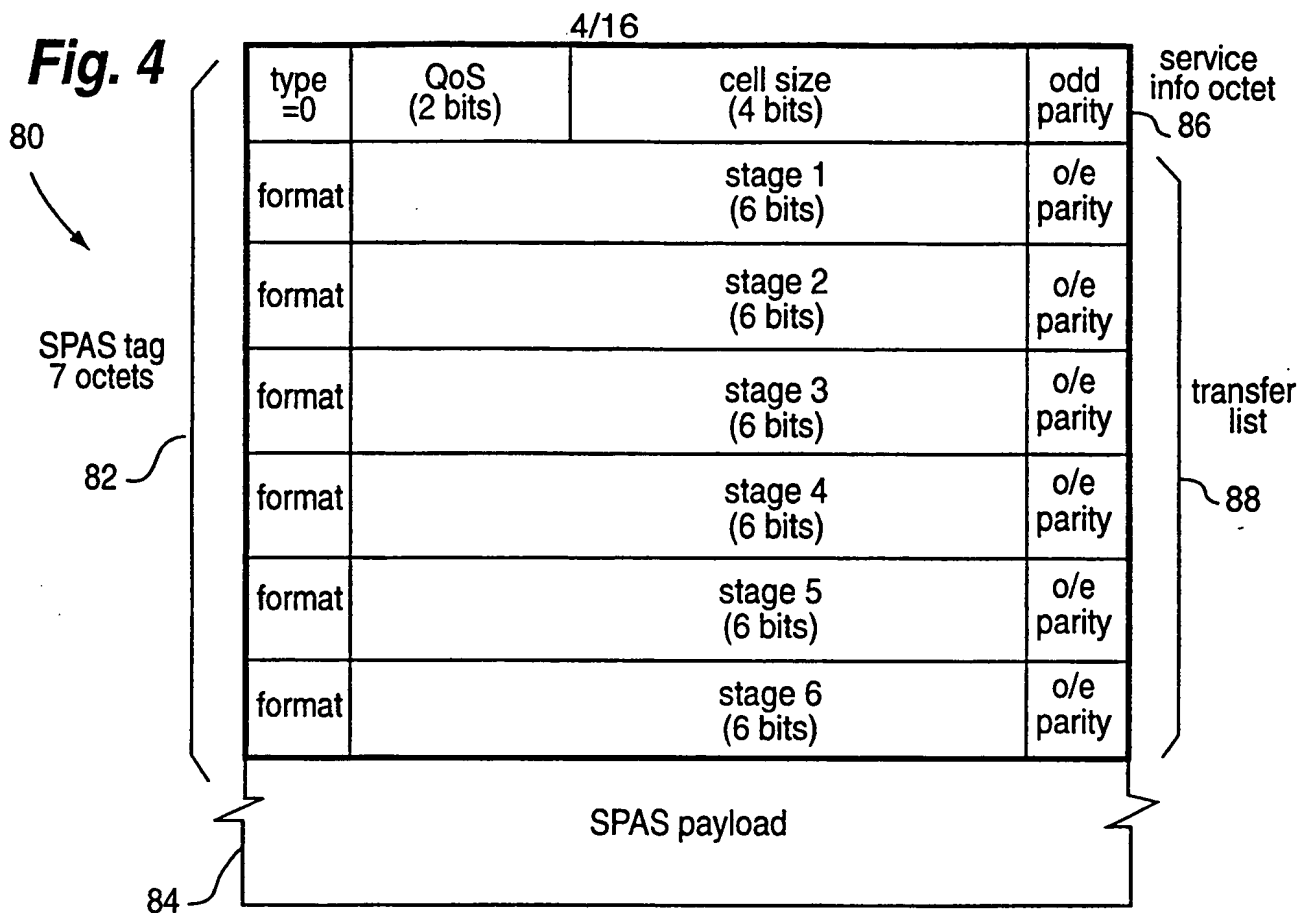
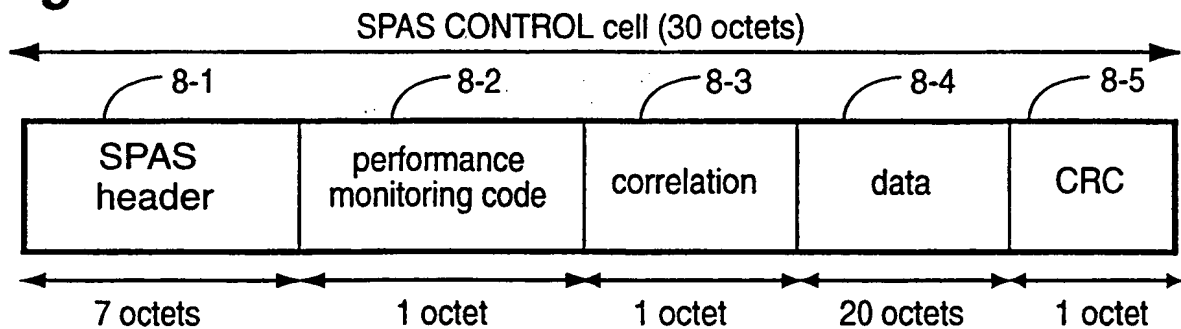



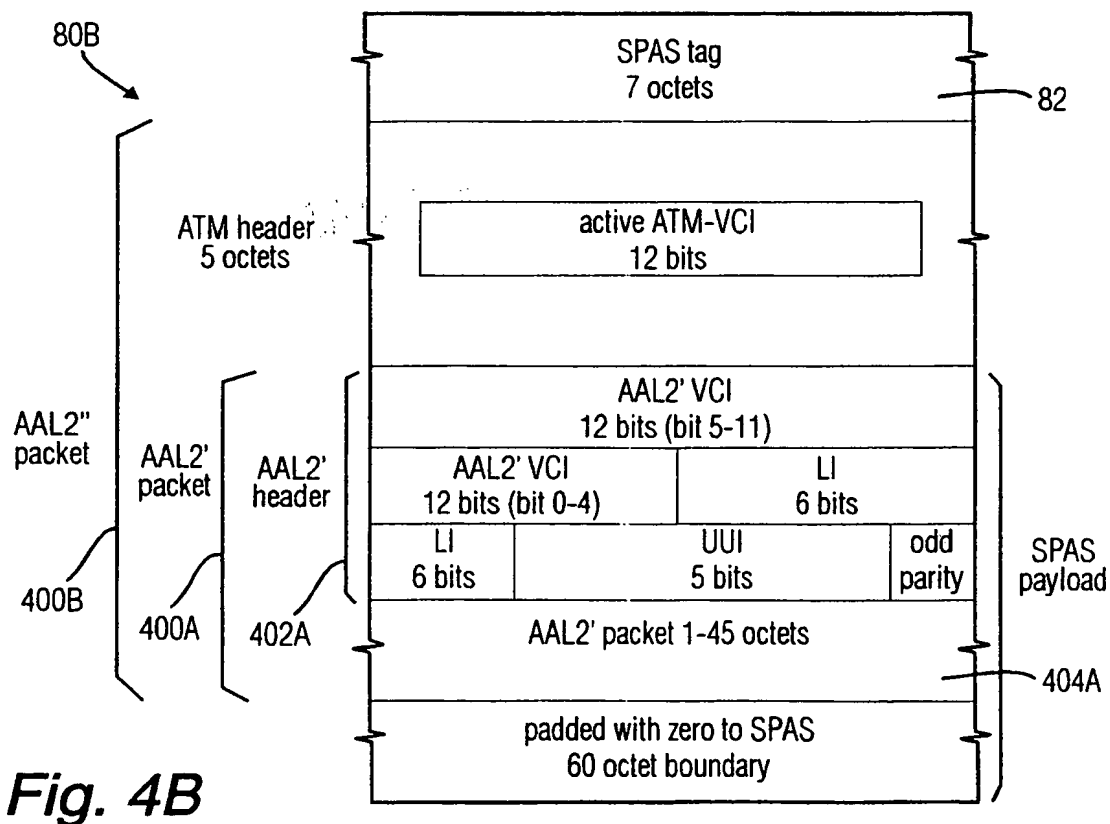
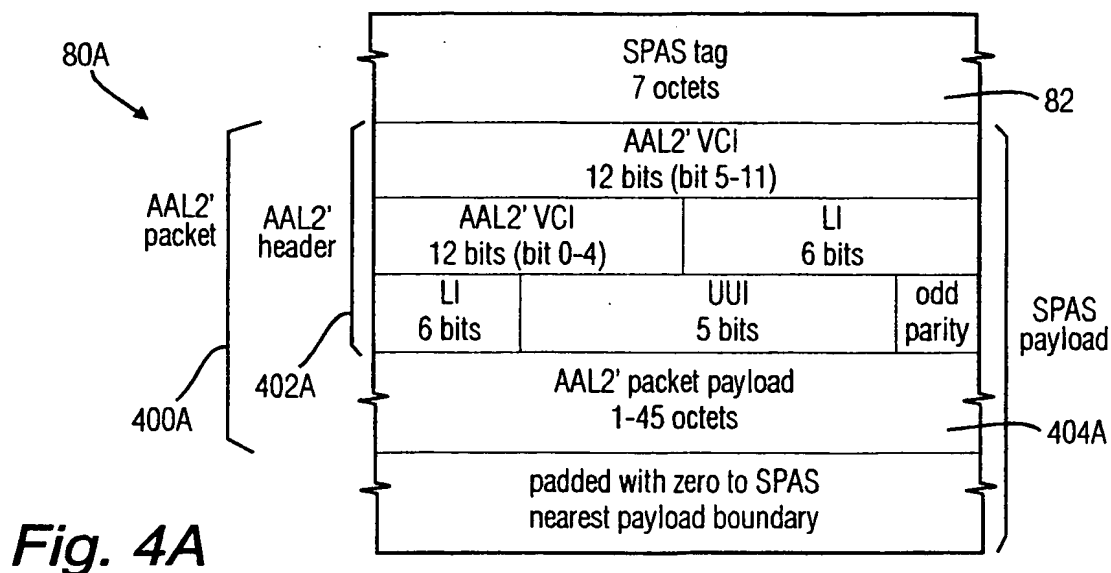
Fig. 3

THIS PAGE BLANK (USPTO)

**Fig. 8****Fig. 9**


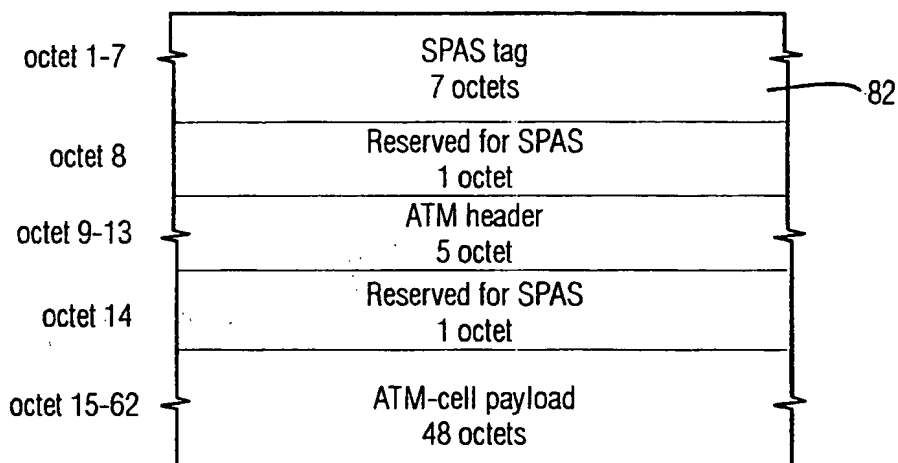
THIS PAGE BLANK (USPTO)

5/16



THIS PAGE BLANK (USPTO)

6/16

*Fig. 4C**Fig. 4D*

THIS PAGE BLANK (USPTO)

7/16




Fig. 5

THIS PAGE BLANK (USPTO)




Fig. 5A




Fig. 5B

THIS PAGE BLANK (USPTO)

9/16




Fig. 5C




Fig. 5E

THIS PAGE BLANK (USPTO)


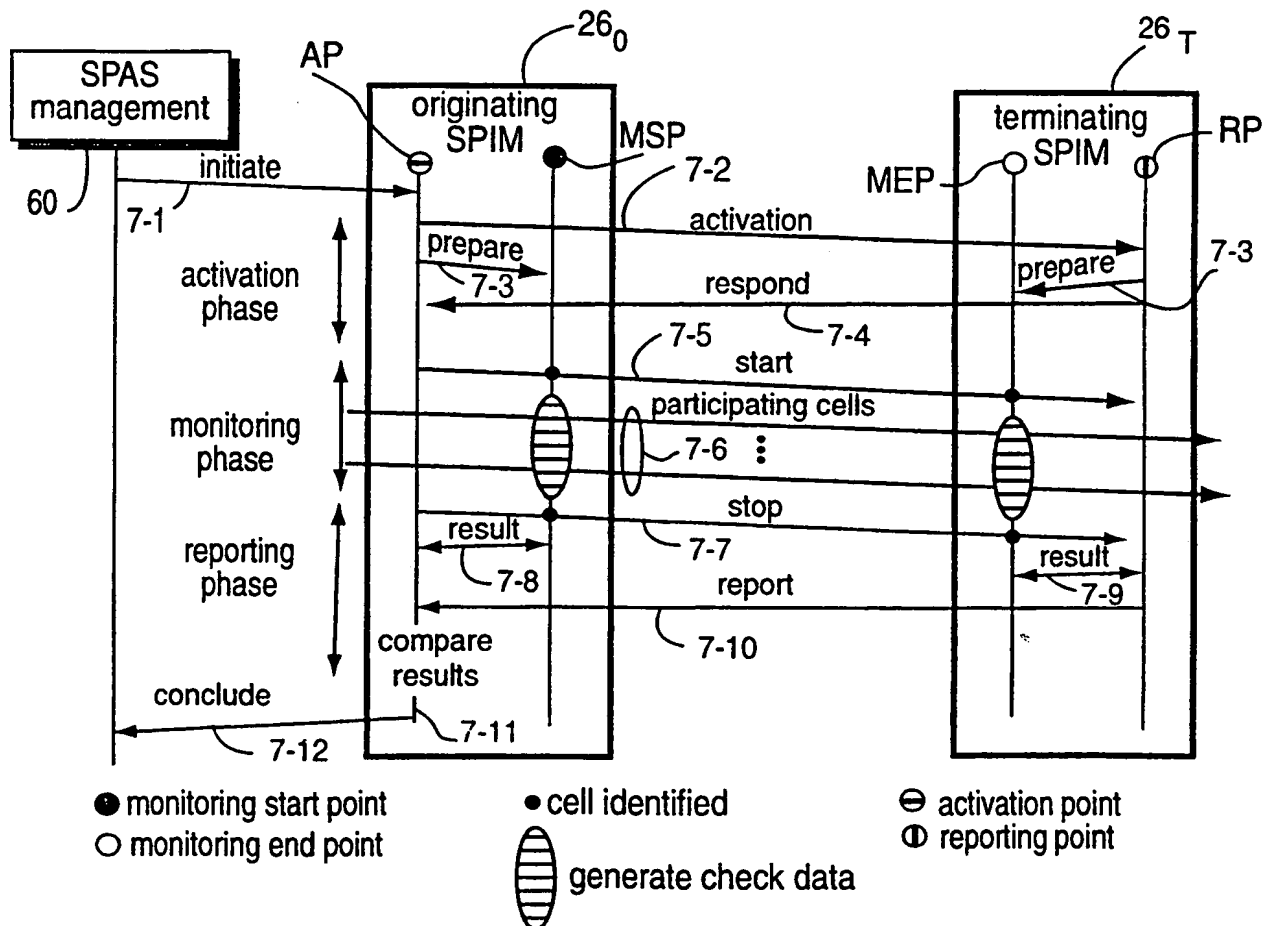



Fig. 5D

THIS PAGE BLANK (USPTO)

11/16

Fig. 6**Fig. 7**

THIS PAGE BLANK (USPTO)




Fig. 10A




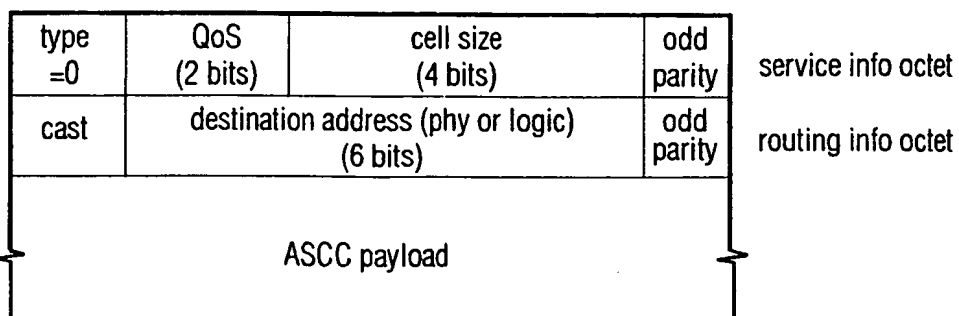



Fig. 10B

THIS PAGE BLANK (USPTO)

13/16

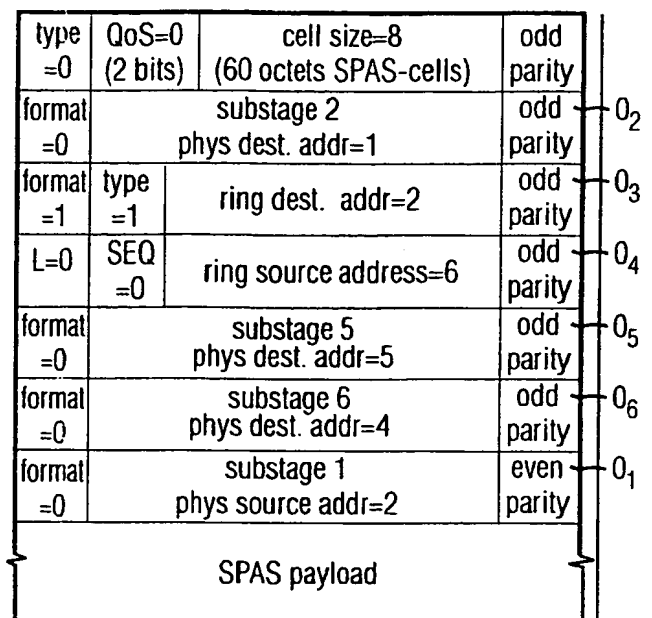
**Fig. 12****Fig. 11****Fig. 15**

THIS PAGE BLANK (USPTO)



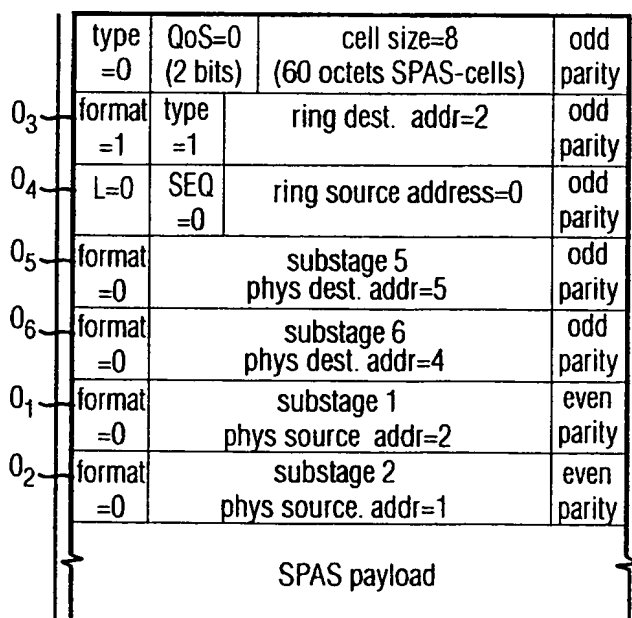
82A

Fig. 13A



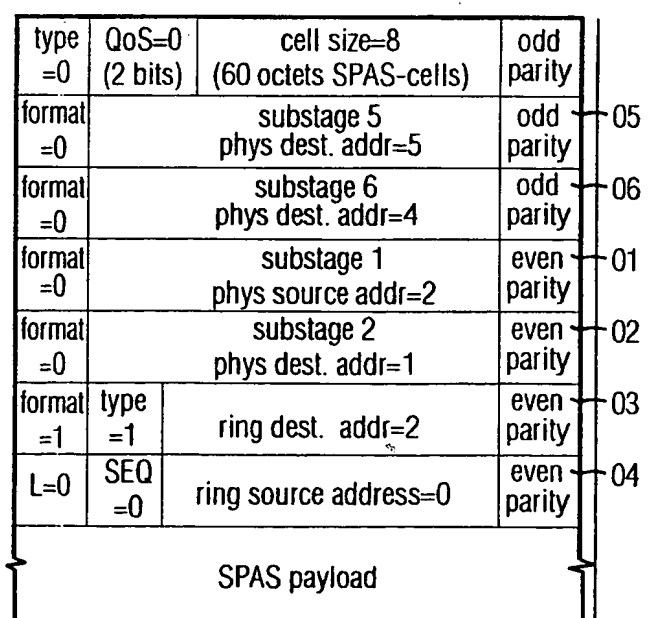
82B

Fig. 13B



82C

Fig. 13C



82D

Fig. 13D

THIS PAGE BLANK (USPTO)

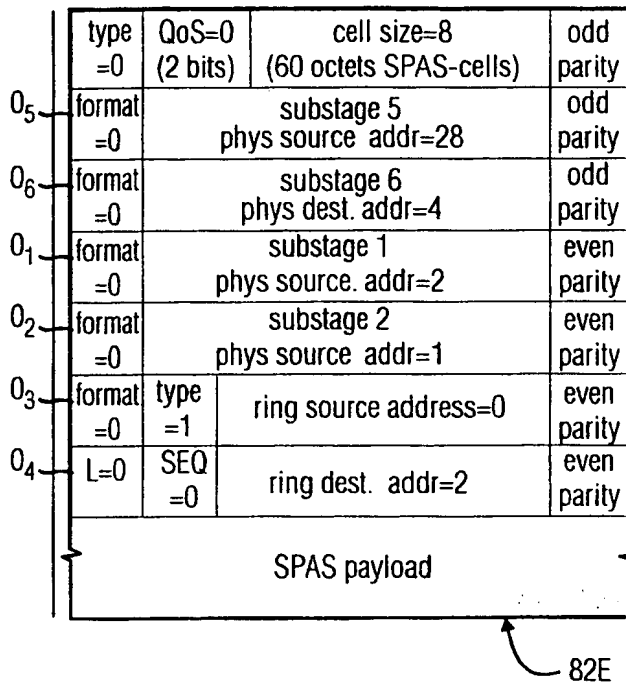


Fig. 13E

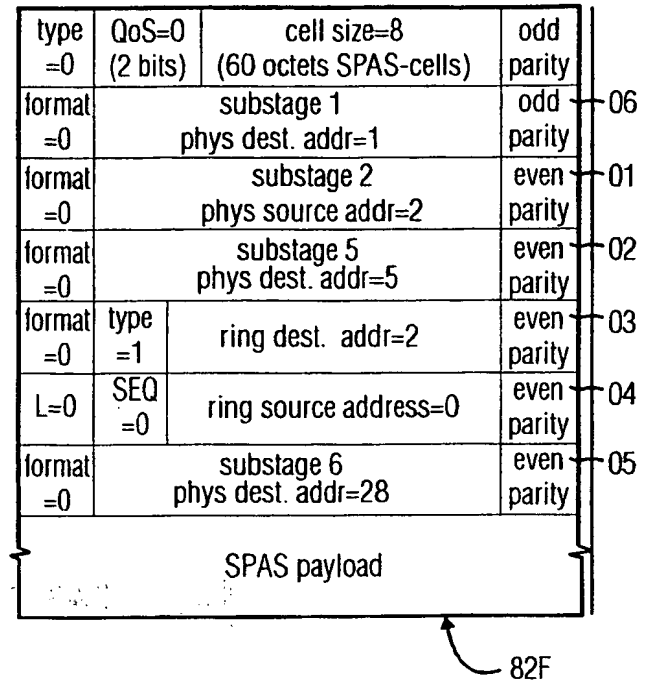
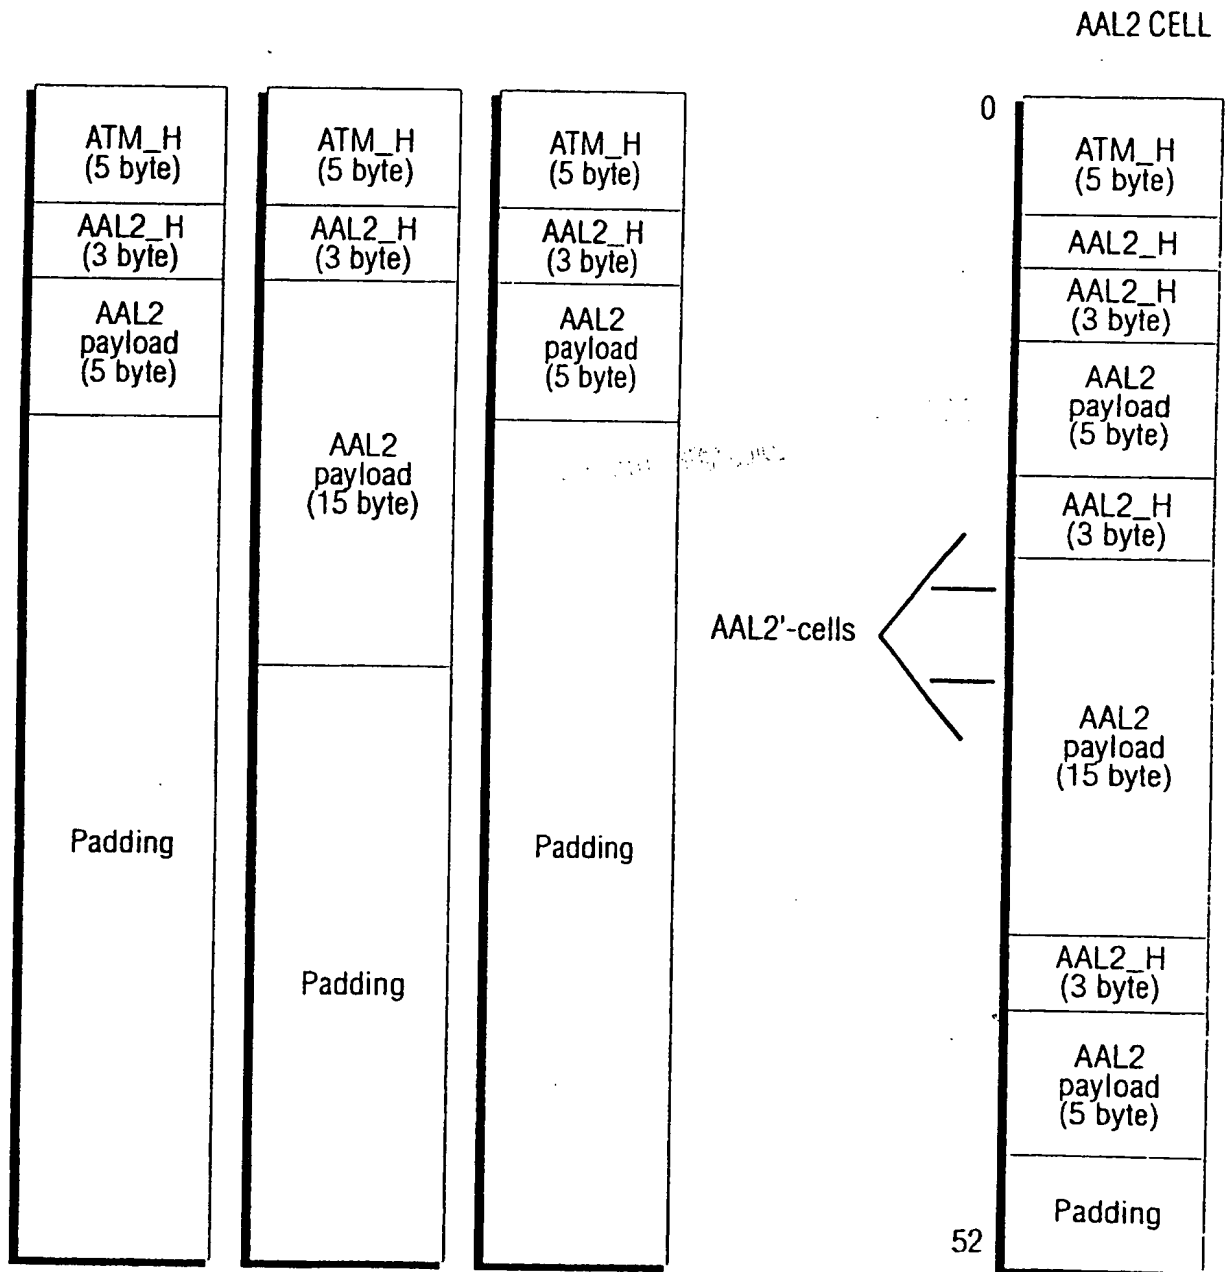


Fig. 13F

THIS PAGE BLANK (USPTO)

16/16

Fig. 14

THIS PAGE BLANK (USPTO)

INTERNATIONAL SEARCH REPORT

International Application No

PCT/SE 99/02319

A. CLASSIFICATION OF SUBJECT MATTER
 IPC 7 H04Q11/04 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04Q H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 441 257 A (HITACHI LTD ;HITACHI SOFTWARE ENG (JP)) 14 August 1991 (1991-08-14) abstract column 2, line 34-38 column 7, line 1-25 column 8, line 56 -column 10, line 14 figure 21	1-3, 7, 10-14, 16-18, 24, 25
X	EP 0 778 719 A (HITACHI LTD) 11 June 1997 (1997-06-11) abstract column 4, line 3-10; figures 1, 3	1, 16
A	US 4 891 802 A (JASMER WOLFGANG E ET AL) 2 January 1990 (1990-01-02) column 7, line 16-22 -/-	7-9

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

2 May 2000

Date of mailing of the international search report

12/05/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
 Fax: (+31-70) 340-3016

Authorized officer

Dhondt, E

THIS PAGE BLANK (USPTO)

INTERNATIONAL SEARCH REPORT

Inte. onal Application No

PCT/SE 99/02319

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	CHIUSSI F M ET AL: "A 0.8- M BICMOS SEA-OF-GATES IMPLEMENTATION OF THE TANDEM BANYAN FAST PACKET SWITCH*" PROCEEDINGS OF THE CUSTOM INTEGRATED CIRCUITS CONFERENCE, US, NEW YORK, IEEE, vol. CONF. 13, page 331-336 XP000295649 ISBN: 0-7803-0015-7 figure 3	7-9
A	GB 2 303 521 A (PMC SIERRA INC) 19 February 1997 (1997-02-19) abstract claim 15	10,12

THIS PAGE BLANK (USPTO)

INTERNATIONAL SEARCH REPORT

Information on patent family members

Initial Application No

PCT/SE 99/02319

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0441257 A	14-08-1991	JP 2834253 B	09-12-1998
		JP 3231540 A	15-10-1991
		CA 2035428 A	08-08-1991
		US 5177736 A	05-01-1993
EP 0778719 A	11-06-1997	EP 0809380 A	26-11-1997
		CA 1291806 A	05-11-1991
		DE 3856370 D	25-11-1999
		EP 0299473 A	18-01-1989
		JP 2561045 B	04-12-1996
		JP 8051435 A	20-02-1996
		JP 2561046 B	04-12-1996
		JP 8051436 A	20-02-1996
		JP 2561047 B	04-12-1996
		JP 8051437 A	20-02-1996
		JP 2561048 B	04-12-1996
		JP 8051438 A	20-02-1996
		JP 2001669 A	05-01-1990
		JP 2569118 B	08-01-1997
		US 4910731 A	20-03-1990
		US 5799014 A	25-08-1998
		US 6016317 A	18-01-2000
		US 5710770 A	20-01-1998
		US RE34305 E	06-07-1993
US 4891802 A	02-01-1990	DE 3714385 A	10-11-1988
		DE 3850440 D	04-08-1994
		EP 0289087 A	02-11-1988
		JP 1016045 A	19-01-1989
		JP 2741864 B	22-04-1998
GB 2303521 A	19-02-1997	CA 2181293 A	18-01-1997
		JP 2965907 B	18-10-1999
		JP 9107369 A	22-04-1997
		US 5889778 A	30-03-1999

THIS PAGE BLANK (USPTO)

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☒ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

THIS PAGE BLANK (USPTO)